arXiv:1702.08342v2 [cs.CR] 28 Feb 2017

Curie* : Policy-based Secure Data Exchange

Z. Berkay Celik[1], *Hidayet Aksu*[2], *Abbas Acar*[2], *Ryan Sheatsley*[1], *A. Selcuk Uluagac*[2], *and Patrick McDaniel*[1]

[1]SIIS Laboratory, Department of CSE, The Pennsylvania State University
{zbc102,mcdaniel,rms5643}@cse.psu.edu
[2]CPS Security Lab, Department of ECE, Florida International University
{haksu,aacar001,suluagac}@fiu.edu

Abstract

Data sharing among partners—users, organizations, companies—is crucial for the advancement of data analytics in many domains. Sharing through secure computation and differential privacy allows these partners to perform private computations on their sensitive data in controlled ways. However, in reality, there exist complex relationships among members. Politics, regulations, interest, trust, data demands and needs are one of the many reasons. Thus, there is a need for a mechanism to meet these conflicting relationships on data sharing. This paper presents Curie, an approach to exchange data among members whose membership has complex relationships. The CPL policy language that allows members to define the specifications of data exchange requirements is introduced. Members (partners) assert who and what to exchange through their local policies and negotiate a global sharing agreement. The agreement is implemented in a multi-party computation that guarantees sharing among members will comply with the policy as negotiated. The use of Curie is validated through an example of a health care application built on recently introduced secure multi-party computation and differential privacy frameworks, and policy and performance trade-offs are explored.

1 Introduction

Inter-organizational data sharing is crucial to the advancement of many domains including security, health care, and finance. Previous works have shown the benefit of data sharing with the use of secure computation and differential privacy guarantees. Secure computation offers data sharing among multiple members while avoiding the risks of disclosing the sensitive data [18]. Recent works in differential privacy allow multiple members to perform joint analysis while protecting the privacy of individuals [33, 44]. These approaches solve the privacy requirements of members' data or individuals, yet do not consider the complex relationships among members on data sharing. Members want to be able to articulate and enforce their conflicting requirements on data sharing.

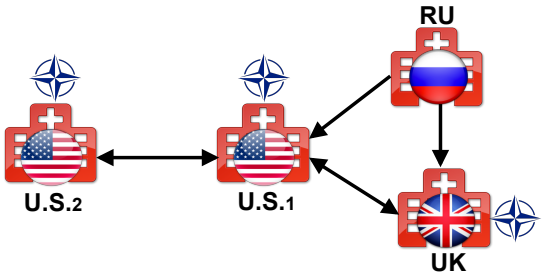

To illustrate complex data sharing requirements, consider health care organizations that collaborate for diagnosis of patients experiencing blood clots. Members wish to dictate their needs through their legal and political limitations as follows. U.S.$_1$ is able to share its complete data for nation-wide members (U.S.$_2$) [35,36], yet it is obliged to share the data of patients deployed in NATO countries with NATO members (UK) [37]. However, U.S.$_1$ wishes to acquire all patient data from other countries. UK is able to share and acquire complete data from NATO members, yet it desires to acquire only data of certain race groups from U.S$_1$ to increase its data diversity. RU wishes to share and acquire complete data from all members, yet members limit their data share to Russian citizens who live in their countries. Such complex data sharing requirements also commonly occur today in non-healthcare systems [29,42]. For instance, National Security Agency has varying restrictions on how human intelligence is shared with other countries; financial companies share data based on trust, and competition among each other.

This paper presents a policy-based data exchange approach, called Curie, that allows data exchange among members that have complex relationships. Members specify their requirements on data exchange using a policy

*Our paper named after Marie Curie. She is physicist and chemist who conducted pioneering research in health care.

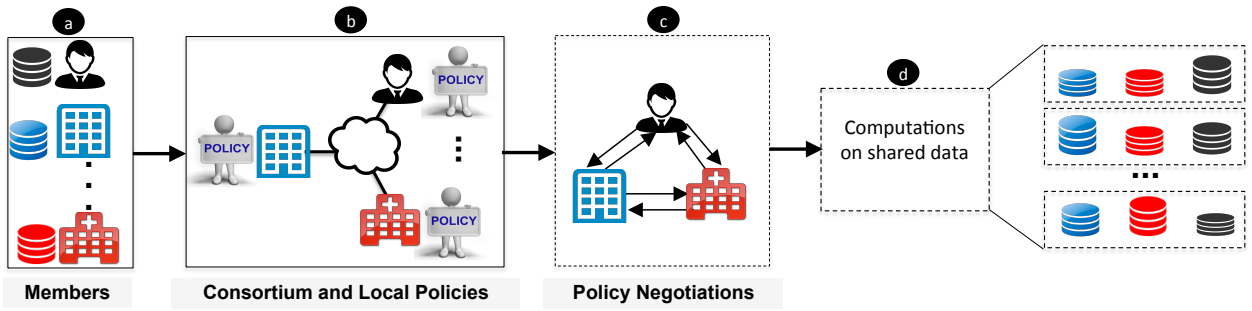

Figure 1: Curie data exchange processes before a computation starts. Dashed boxes indicate data that should remain confidential.

language (CPL). The requirements defined with use of CPL form the local data exchange policies of members. Local policies are defined separately for data sharing and data acquisition policies. This property allows asymmetric relations on data exchange. For example, a member does not necessarily have to acquire the data that the other members dictate to share. By using these two policies, members specify statements of who to share/acquire and what to share/acquire. The statements are defined using conditional and selection expressions. Selections allow members to filter data and limit the data to be exchanged, whereas conditional expressions allow members to define logical statements. Another advanced property of CPL is predefined data-dependent conditionals for calculating the statistical metrics between member's data. For instance, members can define a conditional to compute the intersection size of data columns without disclosing their data. This allows members to define content-dependent conditional data exchange in their policies.

Once members have defined their local policies, they negotiate a global sharing agreement. The guarantee provided by Curie is that all data exchanged among members will respect the agreement. The agreement is compiled in a multi-party privacy-preserving computation enhanced with differential privacy guarantees. In this, we make the following contributions:

- We introduce Curie, an approach for data exchange among members that have complex relationships. Curie includes CPL policy language that allows members to define specifications of data exchange requirements, negotiate a global agreement, and compile the agreements in a multi-party computation such that policies respect the negotiated policy.
- We validate Curie through an example of healthcare application used to prescribe warfarin dosage. A privacy-preserving dose algorithm is compiled with the use of various data exchange policies. We show that Curie incurs low overhead and policies are effective at improving the algorithm outcome.
- We explore a differential privacy mechanism integrated into privacy-preserving dose algorithm and evaluate it through policies. We show that the accuracy improvements obtained from policies vanish for privacy budget $\varepsilon \leq 20$. This result highlights that protecting individual privacy may not be balanced without harming the efficacy of policies.

2 Inter-organizational Data Exchange

Depicted in Figure 1, Curie includes two independent parts: policy management and sharing computation.

Policy Management- We define a *consortium* that is a group made up of two or more members–individuals, companies or governments (ⓐ). Members within a consortium aim to exchange sensitive data toward achieving an objective. For instance, data is curated from medical history of patients or financial reports of companies with the objective to build a predictive model on negotiated data. Data content is an instance of a consortium and may differ from consortium to consortium.

Each member of a consortium defines a *local policy* (ⓑ). The local policy of a member dictates the requirements of data exchange as follows:

- The member wishes to specify with whom to share and acquire data (*partnership requirement*).
- The member wishes to define what data to share and acquire (*sharing and acquisition requirement*).

In this, the member wishes refine its sharing and acquisition requirements to express the following:

- The member wishes to dictate a set of conditions to restrict data sharing and select which data to be acquired (*conditional selective share and acquisition*); and
- The member wishes to dictate conditionals based on the other member's data (*data-dependent conditionals*).

The policy of members need not be–nor are likely to be–symmetric. A local policy is defined with requirements for sharing and acquisition that is tailored to each partner member in the consortium–thus allowing each pairwise sharing to be unique. Here, the local policies are used to negotiate pairwise sharing within the consortium. To

illustrate how members negotiate an agreement, consider the following consortium of three members:

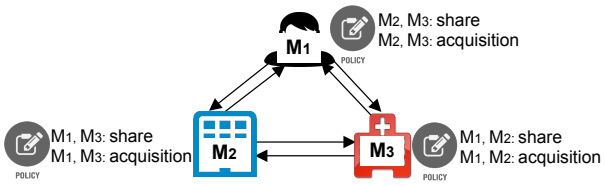

Each member initiates pairwise policy negotiations with other members to reconcile contradictions between acquisition and share policies (ⓒ). A member starts the negotiation by sending a request message including the acquisition policy defined for a member. When a member receives the acquisition policy, it reconciles the received acquisition policy with its share policy specified for that member. Three negotiation outcomes are possible: the acquisition policy is fully satisfied, partially satisfied with the intersection of acquisition and share policies or is an empty set. A member completes its negotiations after all of its acquisition policies for interested parties are negotiated. Each member follows the same process and completes its negotiations. We note that for this process to work securely such negotiation must be done in a way that ensures all policies are correctly shared between partners and that all parties reach a shared consensus on the result. For brevity, we omit the specifics of how these security requirements are achieved, but refer the reader to known multiparty-consensus algorithms [31, 32].

Data Sharing- Once members negotiate their policies (ⓓ), Curie compiles a multiparty data exchange device using secure multi-party computation circuits enhanced with differential privacy guarantees. This device ensures computational and individual privacy. The process on negotiated data is illustrated as follows:

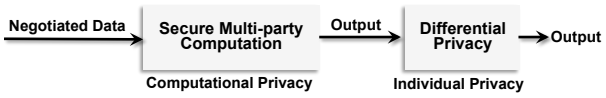

To ensure computational privacy, Curie includes cryptographic primitives of secure multi-party computation that allows members to perform correct computations on negotiated data with no disclosed data from any single member. The primitives are used to implement protocols for secure computations through garbled circuits and homomorphic encryption. To ensure individual privacy, Curie includes a functional mechanism that allows members to apply differential privacy to the output of the secure computation. The functional mechanism ensures that the analysis results containing sensitive information are not traceable back to particular individuals. We use this device for evaluation of Curie through varying consortia and policies.

@M_1	
M_2–	desires to acquire complete data of users who are older than 25
M_2–	shares its complete data
M_3–	desires to acquire Asian users such that the Jaccard similarity of its age column and M_3's age column is greater than 0.3
M_3–	shares its complete data
@M_2	
M_1–	desires to acquire complete data
M_1–	limit its share to EU and NATO citizen users if M_1 is both NATO and EU member and located in North America. Otherwise, it shares only White users
M_3–	desires to acquire complete data if M_3 is a NATO member
M_3–	shares its complete data
@M_3	
M_1–	desires to acquire complete data of users having genotype 'A/A'
M_1–	share complete data if intersection size of its and M_1's genotype column is less than 10. Otherwise, shares data of users that weigh more than 100 pounds
M_2–	desires to acquire complete data
M_2–	shares complete data if M_2 is EU member and its data size is greater than 1K

Table 1: An example of data exchange requirements of members.

3 Policy Description Language

We now illustrate the format and semantics of the Curie policy language (CPL). A BNF description of CPL is presented in Appendix-A. Turning to the example consortium established with three members, each member defines its requirements for other members on a dataset having the columns of age, race, genotype and weight (see Table 1). The criteria defined by members are used throughout to construct their local policies.

Share and Acquisition Clauses- Curie policies are collections of orders clauses. The collection of clauses for partners defines the local policy of a member. The clauses allow each member to dictate an independent consortium-based policy for each other member. Clauses have the following structure:

⟨clause tag⟩ : ⟨members⟩ : ⟨conditionals⟩ :: ⟨selections⟩;

Clause tags are reference names for policy entries. *Share* and *acquire* are two reserved tags. Those clauses are comprised of three parts. The first part, *members*, defines a list of members that specifies with whom to share and acquire. This can be a single member or a comma separated list of members. An empty entry means sharing and acquisition is allowed for all members. The second part, *conditionals*, are a list of conditions controlling when this clause will be executed. A condition is a Boolean function which expresses whether the share or acquire is allowed or not. For instance, a member may define a condition where the data size is greater than a specific value. Only if all

conditions listed in conditionals are true, then this clause is allowed. Last part, *selections*, states what to share or acquire. It can be a list of filters on a member's data. For instance, a member may define a filter on a column of a dataset to limit acquisition to a subset of the dataset. More complex selections can be assigned using member defined sub-clauses. A sub-clause has following structure:

$$\langle \text{tag} \rangle : \langle \text{conditionals} \rangle :: \langle \text{selections} \rangle;$$

where *tag* is the name of sub-clause; conditionals is, as explained above, a list of conditions stating whether this clause will be executed; selections is a list of filters or a reference to a new sub-clause. Complex data selection can be addressed with nested sub-clauses.

CPL allows members to define multiple clauses. For instance, a member may share a distinct subset of data for different conditions. CPL evaluates multiple clauses in a top-down order. When conditionals of a clause evaluate false, it moves to next clause until a clause is matched or it reaches end of the policy file.

Conditionals and Selections- We present the use of conditionals and selections through policies with examples. The format and semantics are detailed. Consider an example of two members, M_1 and M_2, within a consortium. They define their local policies as follows:

$@M_1$

acquire : M_2 : :: s_1 ; (1)

share : M_2 : :: ; (2)

$@M_2$

acquire : M_1 : :: ; (3)

share : M_1 : c_1, c_2 :: fine-select ; (4)

fine-select : c_3 :: s_2 ; (5)

fine-select : :: s_3 ; (6)

where c_1, c_2 and c_3 are conditionals, s_1, s_2 and s_3 are selections and fine-select is a tag defined by M_2.

The acquire clause of M_1 states that data is requested from M_2 after it applies s_1 selection (e.g., age > 25) to its data. In contrast, its share clause allows complete share of its data if M_2 requests. On the other hand, the acquisition clause of M_1 dictates requesting complete data from M_2. However, M_2 allows data sharing if the acquisition clause issued by M_1 holds $c_1 \wedge c_2$ conditions (e.g., is both NATO and EU member). Then, M_2 delegates selection to member-defined fine-select sub-clauses. fine-select states that if the request satisfies the c_3 condition (is located in North America) then the request is met with the data that is selected by the s_2 selection (e.g., limits share of its data to NATO and EU member country citizens). Otherwise, it shares data that is specified by selection s_3 (White users).

CPL supports selections through filters. A filter contains zero or more operations over data inputs describing the share and acquisition criteria to be enforced. Operations are defined as keywords or symbols such as $<$, $>$, $=$, *in*, *like*, and so on. Selections and filters are defined in CPL as follows:

⟨selections⟩ ::= ⟨filters⟩
| ⟨tag⟩

⟨filters⟩ ::= ⟨filter⟩ [‘,’ ⟨filters⟩]

⟨filter⟩ ::= ⟨var⟩ ⟨operation⟩ ⟨value⟩ | ‘’

Selections are executed when conditionals are true. Conditionals can be consortium and dataset specific. For instance, a member may enforce other members to be located in a specific country, to be a part of an alliance such as NATO, and to have data size greater than a particular value. These conditionals do not require any interaction between members, that is, a member decides conditionals without knowing any details about other member's data. However, members may want to define its policies based on the relations between its data and other member's data. To allow such complex share/acquire policies, CPL provides data-dependent conditionals which is used for evaluating a specific, i.e., statistical function over two members data. This computation is secure and does not leak members data.

Data-dependent Conditionals- A member's decision on whether to share or to acquire data can depend on other member's data. Simply put, one example of a data-dependent conditional among two members could be whether the intersection size of the two sets (e.g., a specific column of a dataset) is not too high. Considering such knowledge, a member can make conditional decision about about share or acquisition of that data. For instance, consider a list of confidential IP addresses used for blacklisting the domains. If a member knows that the intersection size is close to zero, then member may dictate an acquire clause to request complete IP addresses from that member [21].

CPL defines a evaluate keyword for data-dependent conditionals through functions on data. Data-dependent conditionals takes the following form:

⟨conditionals⟩ ::= ⟨var⟩‘=’⟨value⟩ [‘,’ ⟨conditionals⟩]
| ‘evaluate’ ‘(’ ⟨data_ref⟩ ‘,’ ⟨alg_arg⟩ ‘,’
⟨threshold_arg⟩ ‘)’ [‘,’ ⟨conditionals⟩] | ‘’

A member using the data-dependent conditionals defines a reference data (data_ref) required for such computation, an algorithm (alg_arg) and a threshold (threshold_arg) that is compared with the output of the computation. CPL includes four algorithms for data-dependent conditionals (see Table 2). To be brief, intersection size measures the size of overlap between two sets; Jaccard index is a statistic measure of similarity between sets; Pearson correlation is a statistical measure of how much two sets are linearly dependent; and Cosine similarity is a measure of similarity between two vectors. Each algorithm is

Algorithm	Output	Private protocol	Proof
Intersection size	$\lvert\mathcal{D}_i \cap \mathcal{D}_j\rvert$	Set intersection cardinality	[15]
Jaccard index	$\frac{\lvert\mathcal{D}_i \cap D_j\rvert}{\lvert\mathcal{D}_i \cup \mathcal{D}_j\rvert}$	Jaccard similarity	[3]
Pearson correlation	$\frac{COV(\mathcal{D}_i, \mathcal{D}_j)}{\sigma_{\mathcal{D}_i}\sigma_{\mathcal{D}_j}}$	Garbled circuits	[26]
Cosine similarity	$\frac{\mathcal{D}_i\mathcal{D}_j}{\lVert\mathcal{D}_i\rVert\lVert\mathcal{D}_j\rVert}$	Garbled circuits	[26]

Table 2: Algorithms for data-dependent conditionals ($\mathcal{D}_i$ and D_j are the inputs, and μ is for mean and σ is for std. deviation).

based on a different assumption about the underlying reference data. However, central to all of them is to privately (without leaking any sensitive data) measure a relation between two members' data to offer an effective data exchange. We note that these algorithms are found to be effective in capturing input relations in datasets [21, 22].

Data-dependent conditionals are implemented through private protocols (as defined in Table 2). These protocols are implemented with the cryptographic tools of garbled circuits and private functions. Protocols preserve the confidentiality of data. That is, each member gets the output indicated in Table 2 without revealing their sensitive data in plain text. After private protocol terminates, a member compares the output of the algorithm with a threshold value set by the requester. If the output is below the threshold value, the conditional is set true.

Turning to above example M_3 joins the consortium. M_1 and M_2 extend their local policies for M_3 as follows:

$@\mathsf{M}_1$

acquire : M_3 : evaluate(local data, (7)
'Jaccard', 0.3) :: race=Asian;

share : M_3 : :: ; (8)

$@\mathsf{M}_2$

acquire : M_3 : M_3 in $NATO :: ; (9)

share : M_3 : :: ; (10)

$@\mathsf{M}_3$

acquire : M_1 : :: Genotype = 'A/A' ; (11)

share : M_1 : evaluate(local data, (12)
'intersection size', 10) :: ;

share : M_1 : :: weight>150 ; (13)

acquire : M_2 : :: ; (14)

share : M_2 : M_2 in $EU, size(data)> 1K :: ; (15)

The acquire clause of M_1 defines a data-dependent conditional for M_3. M_1 specifies its local data, defines Jaccard index as an algorithm and sets the threshold value equal to 0.3. M_3 evaluates the function, and if its output threshold holds, it agrees to share data if its conditionals evaluate(...) in its share clause defined for M_1 returns true. Otherwise, it consults the next share clause defined for M_1 which states that an individual's weight greater than 150 pounds will be shared. All other share and acquire clauses are trivial. Members agree to share and acquire complete data based on data size (data size > 1K) and alliance membership (NATO member) conditionals.

Putting the pieces together, CPL allows members independently to define a data exchange policy with share and acquire clauses. The policies are dictated through conditionals and selections. This allows members to dictate policies in complex and asymmetric relationships. Defined in Section 2, CPL provides members to dictate the requirements of partnership, share, acquisition, and data-dependent conditionals.

Policy Negotiation- Data exchange between members is governed by matching share and acquire clauses in each members respective policies. Both share and acquire clauses state conditions and selections on the data exchanged. Consider two example local policies with a share clause $@m_2$ ($share : m_1 : c_1 :: s_1$) and matching acquire clause $@m_1$ ($acquire : m_2 : c_2 : s_2$). Curie's negotiation algorithm respects both autonomy of the data owner and the needs of the requester. Consequently, it conservatively negotiates the union of (logical AND) conditionals and selections in resulting policy assignment. The resulting resolved policy in this example is $share : m_1 : c_1 \wedge c_2 :: s_1 \wedge s_2$ which states that the data exchange from m_2 to m_1 is subject to both c_1 and c_2 conditionals and resulting sharing will have s_1 and s_2 selections on m_2's data. This authoritative negotiation makes sure no member's data is shared beyond its explicit intent, regardless how the other members' policies are defines.

Our investigation shows that the negotiation of policies is tractable [31]. This is because negotiation fulfilling the criteria for each clause is based on the union of logical expressions defined in two policies. Each member runs the negotiation algorithm for members found in their member list. After all members terminate their negotiations, the negotiated data will be used in computations.

4 Computations on Negotiated Data

This section presents mechanisms used for secure computation enhanced with differential privacy that allows members to carry out computations over negotiated data. The introduced mechanisms are used for evaluation of Curie in the following section. The guarantee provided by Curie in these mechanisms is that all computations respect the policy as negotiated. As techniques developed within Curie are agnostic to the computations performed over negotiated data, we also present integration of Curie into off-the-shelf frameworks.

Secure Multi-Party Computation- The secure computation ensures that each member of a consortium learns nothing about their data, but the correct computation out-

put is returned. After the introduction of secure two-party computation [46], many practical implementations of secure multi-party computation (SMC) have been proposed [5, 10, 12]. Since then, more efficient protocols and implementations have been proposed to compute a broad class of protocols in a semi-honest threat model [29]. In broader terms, the approaches used for secure computation, which include garbled circuits [2, 46], homomorphic encryption (HE) [13, 23] and oblivious transfer [34], have shown the promise of producing efficient protocols. For evaluation purposes, Curie implements HE because of its computational efficiency. We found that the communication complexity of HE did not require any interactions among members other than the computations.

HE is an encryption system with special properties. It is characterized by four functions [23]: key generation, encryption, decryption, and evaluation. Unlike other functions found in standard cryptographic systems, evaluation is a HE-specific algorithm. It takes ciphertexts as input and computes generic operations on the encrypted data. For instance, an additive homomorphic encryption scheme allows a third party to compute $E(m_1 + m_2)$ by using $E(m_1)$ and $E(m_2)$ without knowing the plaintexts of m_1 and m_2, where E is the encryption function. Curie integrates an open-source Homomorphic Encryption (HElib) software library to implement these functions [38].

Differential Privacy (DP)- DP is a mechanism for releasing statistical results in such a way that the released analysis, which is computed from sensitive information, is not traceable back to particular individuals. Formally, given a summary statistic called $\mathcal{A}()$, an adversary proposes two dataset $\mathcal{S}$ and $\mathcal{S}'$ that differ by only one row, and a test set Q. $\mathcal{A}()$ is called ε-differentially private iff [17]:

$$Pr[\mathcal{A}(\mathcal{S}) \in Q] \leq exp(\varepsilon) \times Pr[\mathcal{A}(\mathcal{S}') \in Q] \qquad (1)$$

where the parameter ε is the privacy budget that controls the trade-off between the accuracy of the differentially-private Q and the amount of information it leaks.

Curie includes a functional mechanism to construct differentially-private algorithms (i.e., regression) over negotiated data [45, 47]. This technique accepts dataset ($\mathcal{D}$), objective function ($f_{\mathcal{D}}(\eta)$), and privacy budget (ε) as an input and returns ε-differentially-private coefficients $\widetilde{\eta}$. The intuition behind the functional mechanism is perturbing the objective function of the optimization problem. Perturbation includes both sensitivity analysis and Laplacian noise insertion as opposed to perturbing the results through differentially-private synthetic data generation.

Integration into off-the-shelf Frameworks- One important advantage of Curie is that it can easily be integrated with existing frameworks due to its modular design. In essence, Curie manages complex relations via policies expressed in CPL and negotiates them. Negotiated data can be used by any computation framework for different objectives. There exist various secure computation and differential privacy frameworks. For instance, our investigation shows that Curie can easily be integrated into ObliVM which is a framework used for automated secure data analysis [30], and others such as RMind [4] and SPARK [18] which are designed for statistical analysis in health care. Concisely, Curie can be used as a middleware between members and such frameworks as a data exchange policy layer.

Input	Values	Category	Description
Age	Years	Demographic	Binned age reported in years.
Height	cm	Background	—
Weight	kg	Background	—
VKORC1	A/A, A/G, G/G	Genotype	Genotype at the -1639 A>G SNP in the VKORC1 promoter.
CYP2C9	*1/*1, *1/*2, *1/*3, *2/*2, *2/*3, *3/*3	Genotype	The alleles are used to reliably estimate their contribution to Warfarin dose.
Race	Asian, Black or African American, Caucasian or White	Demographic	As defined by the U.S. Office of Management and Budget
Enzyme Inducer	Yes or No	Background	Inducers considered are rifampin or rifampicin (Rifadin, Rimactane), phenytoin (Dilantin), and carbamazepine (Tegretol).
Amiodarone	Yes or No	Background	Medication used to treat and prevent irregular heartbeats

Table 3: Patient inputs collected independently by members. The inputs are used for the purpose of accurately estimating the personalized warfarin dose.

5 Example Application: Warfarin Dosing

In order to discuss how Curie works in a realistic setting, we illustrate an actual use of Curie for the prescription of the warfarin medicine. Medical institutions locally collect patient data for the purpose of building an algorithm to estimate the proper personalized dose of warfarin. In this section, we first briefly introduce warfarin, its pharmacogenetic algorithm, members involved in the data collection, secure and differentially-private dose algorithm.

5.1 Warfarin

Warfarin, known as the brand name Coumadin, is a widely (over 20 million times each year in the United States) prescribed anticoagulant medication. It is mainly used to treat (or prevent) blood clots (thrombosis) in veins or arteries. Current clinical practices suggest a fixed initial dose of 5

Policy id	Consortium name	Definition	Acquisition policy?	Share policy?
P.1	Single Source	Each member uses its local patients to build warfarin dose algorithm.	✗	✗
P.2	Nation-wide	Members in the same country establish a consortium based on state and country laws.	✓	✓
P.3	Regional	Members in the same continent establish a consortium.	✓	✓
P.4	NATO-EU	NATO and EU members establish a consortium independently based on mutual agreements.	✓	✓
P.5	Global	Members exchange their complete data to build the dose algorithm.	✓	✓

Table 4: Consortia and data exchange policies constructed among members.

or 10 mg/day. Patients regularly have a blood test to check how long it takes for blood to clot (international normalized ratio (INR)). The normal range of INR for a healthy person is 0.8-1.2, and patients taking warfarin have 2.5 to 3.5 units of INR. Based on the INR, subsequent doses are adjusted to maintain the patient's INR within the desired level. Therefore, it is important to prescribe proper dosages for the patients.

Pharmacogenetic Algorithm for Warfarin Dose- To determine the proper personalized warfarin dosage, a long line of work concluded with an algorithm of an ordinary linear regression [27]. In ordinary least-squares regression, the model is a function $f : \mathcal{X} \to \mathcal{Y}$ that aim at predicting targets of warfarin dose $y \in \mathcal{Y}$ given explanatory inputs of patients $x \in \mathcal{X}$. We represent the dataset of each member $\mathcal{D}_i = \{(x_i, y_i)\}_{i=1}^n$, and a loss function $\ell : \mathcal{Y} \times \mathcal{Y} \to [0, \infty)$. The loss function penalizes deviations between true dose and predictions. Learning is then searching for a dose algorithm f minimizing the average loss:

$$\mathcal{L}(\mathcal{D}, f) = \frac{1}{n} \sum_{i=1}^{n} \ell(f(x_i), y_i). \tag{2}$$

Pharmacogenetic algorithm reduces to minimizing the average loss $\mathcal{L}(\mathcal{D}, f)$ with respect to the parameters of the model f. The algorithm is linear, i.e., $f(x) = \alpha^\top x + \beta$, and the loss function is the squared loss $\ell(f(x), y) = (f(x) - y)^2$. The dosing algorithm gives as well or better results than other more complex numerical methods and outperforms fixed-dose approach[1] [27]. We re-implement the algorithm in Python by direct translation from the authors' implementation and find that the accuracy of our implementation has no statistically significant difference.

5.2 Deployment Environment

This section introduces members in which Curie is deployed. Requirements and policies appropriate for these environments are described. We note that policy construction is a subjective enterprise. Depending on the nature and constraints of a given environment, any number of policies are appropriate. Such is the promise of policy defined behavior; alternate interpretations leading to other application requirements can be addressed through policy.

Members- 24 medical institutions from nine countries and four continents curated the largest and most comprehensive study for evaluating personalized warfarin dose (see Appendix B for details of members involved in the study). They are located in Taiwan, Japan, Korea, Singapore, Sweden, Israel, Brazil, United Kingdom, and the United States. Members collect 68 inputs from patients' genotypic, demographic, background information, yet a long study concluded that eight inputs are sufficient for proper prescriptions (see Table 3 for input definitions). The dataset used in Curie experiments contains 5700 patient records from 21 members.

Illustrating Policies-We define consortia among medical institutions that members state partnerships for data exchange. Table 4 summarizes the consortia. The consortia are defined based on statute and regulations between members, as well as regional, and national partnerships are studied [35–37, 39]. For example, NATO allied medical support doctrine allows strategic relationships that otherwise not obtainable by non-NATO members. Each member in a consortium exchanges data with other member based on its local policy. Various acquisition and share policies through conditionals and selections are studied.

5.3 Secure Dose Algorithm

In this section, we build a secure dose algorithm through policies dictated by members. Members define complete patient inputs as confidential. They build a dose algorithm defined by the output of policy negotiations.

To construct the secure algorithm, we define the formula stated in Equation 2 in a matrix form by minimizing with maximum likelihood estimation:

$$\beta = (\mathcal{X}^\intercal \mathcal{X})^{-1} \mathcal{X}^\intercal \mathcal{Y} \tag{3}$$

where $\mathcal{X}$ is the input matrix, $\mathcal{Y}$ is the dose matrix and β is the algorithm coefficients.

The computation of *local dose algorithm* is straighforward: a member obtains the algorithm coefficients through Equation 3 with use of locally collected patient data. To

[1]The algorithm has been released online http://www.warfarindosing.org to help doctors and other clinicians for predicting ideal dose of warfarin.

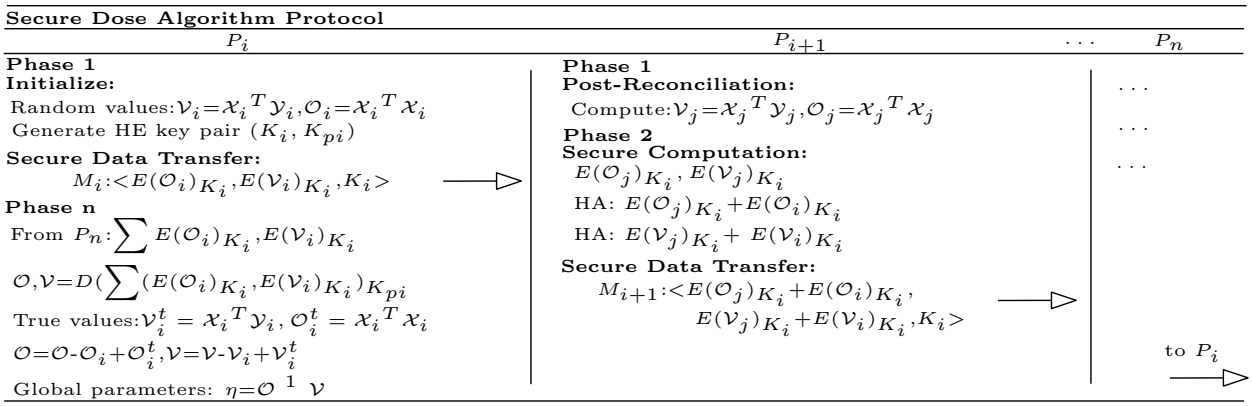

Figure 2: Secure dose algorithm protocol: Member (P_i) starts the protocol, the procedures and message flow of members are highlighted in boldface. At the final phase, P_i is able to compute the dose algorithm parameters of negotiated data.

implement a secure algorithm among consortia members, members translate the data exchanged into neutral input matrices [43]. More particularly, patient samples to be exchanged by each member are computed as an input matrix $\mathcal{X}_0, \ldots, \mathcal{X}_n$ and dose matrix $\mathcal{Y}_0, \ldots, \mathcal{Y}_n$. The transformation defines each members' *local statistics* $\mathcal{O}_i = \mathcal{X}^\intercal \mathcal{X}$ and $\mathcal{V}_i = \mathcal{X}^\intercal \mathcal{Y}$. Local statistics define the output of the negotiation defined for each member in a consortium. The aggregation of these subsets corresponds to a *pooled dataset* which is the exact amount a member negotiates to obtain from other members. Dose algorithm construction from the pooled dataset is defined as concatenation of local statistics as follows:

$$\mathcal{X}^\intercal \mathcal{X} = \left[\mathcal{X}_1^\intercal | \ldots | \mathcal{X}_n^\intercal\right] \left[\mathcal{X}_1 | \ldots | \mathcal{X}_n\right]^\intercal = \sum_{i=1}^{n} \mathcal{X}_i^\intercal \mathcal{X}_i = \sum_{i=1}^{n} \mathcal{V}_i = \mathcal{V}$$
$$\mathcal{X}^\intercal \mathcal{Y} = \left[\mathcal{X}_1^\intercal | \ldots | \mathcal{X}_n^\intercal\right] \left[\mathcal{Y}_1 | \ldots | \mathcal{X}_n\right]^\intercal = \sum_{i=1}^{n} \mathcal{X}_i^\intercal \mathcal{Y}_i = \sum_{i=1}^{n} \mathcal{O}_i = \mathcal{O} \quad (4)$$

A member computes the parameters of the algorithm using sum of other members local statistics with Equation 4. The local statistics includes $m \times m$ constant matrices where m is the number inputs (independent of number of patient samples). Using this observation, a party computes the global coefficients as follows:

$$\eta^{(global)} = (\mathcal{X}^\intercal \mathcal{X})^{-1} \mathcal{X}^\intercal \mathcal{Y} = \mathcal{O}^{-1} \mathcal{V} \quad (5)$$

In Equation 5, while the accuracy objective of secure computation is guaranteed using the coefficients obtained from sum of local statistics, exchange of clear statistics among parties may leak information about members' data. A member is able to infer knowledge about the distribution of each input of other members from matrices of $\mathcal{O}_i$ and $\mathcal{V}_i$ [18]. Furthermore, an adversary may sniff data traffic to control and modify exchanged messages (We detail the security evaluation in Appendix D).

To solve these problems, we use homomorphic encryption that allows computation on ciphertexts. Depicted in Figure 2, an example session of n members is authorized for data exchange in a consortium. In this example, a ring topology is used for secure group communication (i.e., P_i talks to P_{i+1}, and similarly P_n talks to P_i). P_1 initially generates a pair of encryption keys using the homomorphic cryptosystem and broadcasts the public key to the members in its member list. P_1 then generates random $\mathcal{V}_i$, $\mathcal{O}_i$ and encrypts them $E(\mathcal{O}_i)_{K_i}$ and $E(\mathcal{V}_i)_{K_i}$ using its public key K_i. It starts the session by sending them to the next member in the ring. When next member receives the encrypted message, it adds its local $\mathcal{V}_i$ and $\mathcal{O}_i$ matrices through homomorphic addition to the output of its reconciliation algorithm for P_1 and passes to the next member. Remaining members take the similar steps.

At the last stage of the protocol, P_1 receives the sum statistics of $\mathcal{O}_i$ and $\mathcal{V}_i$ from P_n. It subtracts the initial random values of $\mathcal{V}_i$, $\mathcal{O}_i$ and adds its true values used for dose algorithm. In such protocol, all data and message exchanges are encrypted, and a member is only able to infer the global dose algorithm and that of without its contribution (see Appendix D for details).

5.4 Differentially-Private Dose Algorithm

In the last section, we saw that members build the secure dose algorithm on negotiated data through their policies. In this section, we consider individual privacy that guarantees no information leakage on the targeted individual involved in the computation with high confidence (controlled by privacy budget). Specifically, while members compute a secure algorithm using the data obtained as a result of their negotiations, they also ensure that an adversary cannot infer whether an individual is included in computations to learn the algorithm. In warfarin study, this corresponds to differentially-private secure dose algorithm on negotiated data.

To build differentially-private secure algorithms, we

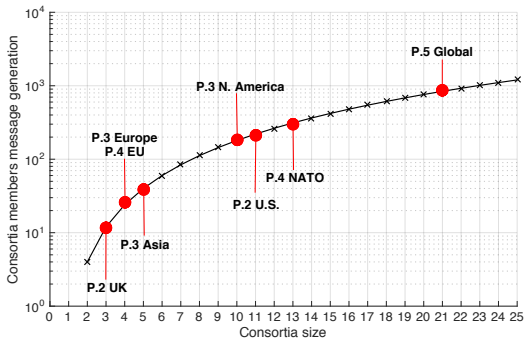

Figure 3: Policy negotiation cost - Costs associated with number of varying members in a consortium.

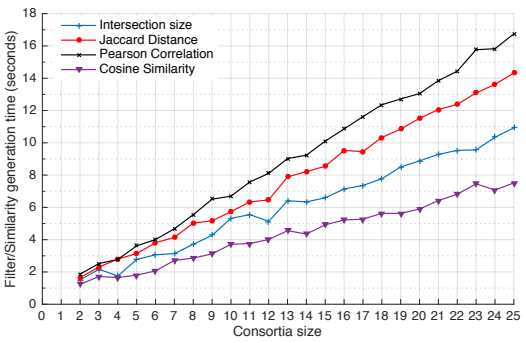

Figure 4: Selections and Data-dependent conditionals costs - Costs associated with varying members and algorithms.

use the functional mechanism introduced in Section 4. To inter-operate the functional mechanism with the secure algorithm, members convert each column from [min, max] to [-1,1] before negotiation starts. This processing ensures that sufficient noise is added to the objective function on negotiated data. Then members proceed with the protocol. At the final stage of the secure algorithm protocol, a member gets clear statistics of $\mathcal{O} = \mathcal{X}^\intercal \mathcal{X}$ and $\mathcal{V} = \mathcal{X}^\intercal \mathcal{Y}$ and input dimension d by means of size of $\mathcal{O}$ or $\mathcal{V}$. These statistics are the exact quantities that is intrinsically minimized in the objective function of the functional mechanism [47]. Having these inputs, a member may compute ε-differential private secure algorithm that does not require any additional data exchange and computational overhead.

6 Evaluation

This section details the operation of the Curie through policies. We show how flexible data exchange policies are implemented and operated through Curie framework. We focus on the following questions:

1. Can members reliably use the Curie to implement application-specific policies?
2. Does Curie support data exchange requirements for multiple consortia?
3. Are there performance trade-offs in configuring data exchange policies?
4. Does establishing a consortium through policies have an impact on application results?

The answers to the first three question are addressed in Section 6.1, and the last question is answered in Section 6.2. As detailed next, Curie incurs low computational overhead across a large number of members and data samples, oftentimes less than a minute. We show how flexible data exchange policies can improve—oftentimes substantially—the accuracy, and minimizes the adverse outcomes. We evaluate differential-private algorithm through varying policies. We found that preserving the individual privacy may not be balanced without harming the efficacy of policies (See Section 6.3).

6.1 Performance

We present the findings of warfarin study in terms of the costs associated with the various Curie mechanisms. We illustrate the cost of the data exchange policies and secure computation; policy negotiation, data-dependent conditionals, and secure dosing algorithm.

The experiments described in this section were performed on a cluster of machines with 32 GB of maximum memory and 16-core Intel Xeon CPU at 1.90 GHz running Red Hat Enterprise Linux Server, where we utilized only one core.

The Curie secure computation protocols are implemented using the open-source Homomorphic Encryption (HElib) library [38]. Our implementation uses Brakerski-Gentry-Vaikuntanathan (BGV) fully homomorphic cryptosystem [7]. This setup consists of many optimizations to make the homomorphic evaluation runs faster with effective use of the Smart-Vercauteren ciphertext packing techniques and Gentry-Halevi-Smart optimizations [38]. We calibrate level and slot parameters of BGV for encrypting multiple messages at one time through SIMD (Single-Instruction Multiple-Data) techniques and set the security parameter as 128 bits. These parameters are optimized per member to increase the number of allowed homomorphic operations without decryption failure and to reduce the computational time.

The first set of experiments sought to characterize the policy construction and policy negotiation costs. Consortia are instrumented to classify overheads incurred by a number of messages and time required to calculate selections and data-dependent conditionals. All the costs not specific to the policies are not included in test measurements (e.g., network latency). The benchmark results are summarized in Figure 3 and 4 and discussed below.

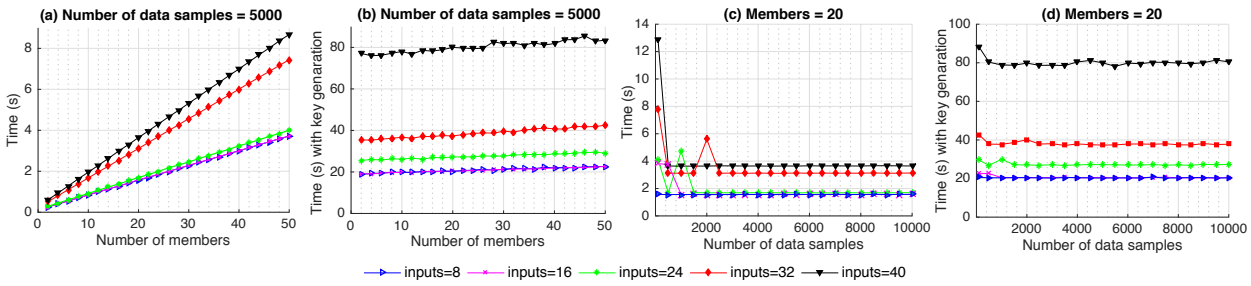

Figure 5: Secure algorithm performance - Std. dev. of ten runs is $\pm$ 3.6 and $\pm$ 0.3 (s) with and without homomorphic key generation.

Figure 3 shows the number of messages for policy construction required for different consortia size. The number of members in a consortium is labeled. For instance, NATO consortium has 13 members; ten members from U.S. and three from the UK. Since the data exchange relations are not symmetric, each member sends acquisition request to its partners and gets the output of negotiation (not raw data). We note that costs associated with selection and data-dependent conditionals dictated by the members do not require additional messages. The acquisition request includes variables when conditionals are defined (e.g., reference data for data-dependent conditionals), and the result is returned with the negotiation output message. However, the use of the selections and data-dependent conditionals brings additional processing cost.

In Figure 4, we show the costs associated with the use of selection and data-dependent conditionals. Data-dependent conditionals and selections are dictated by all the members. The input size for the data-dependent constraint computation is set to 200 real values. This is the average number of inputs found in members' database. Since selections do not require a computational overhead, it yields a processing time of milliseconds. However, the computation of data-dependent constraint metrics requires the use of the specific secure functions. The costs associated with varying data-dependent constraint metrics depends on the secure protocol of associated function. In our experiments, cosine similarity and intersection size exhibited shorter computation time than Pearson correlation and Jaccard distance. While implementation dependent, 25 members are able to compute metrics less than 18 seconds. Note that the results serve as an upper bound; all the members are assumed to define data-dependent conditionals. Future work will include an investigation of network latency and throughput within increasing consortia sizes.

Our second series of experiments characterizes the effects of policy construction on secure computation. These experiments measure the average time of secure dose algorithm with varying number of members and data samples. Evaluations are repeated with input size of 8, 16, 24, 32, and 40. The results are summarized in Figure 5.

Computations are instrumented to classify the overheads incurred by the cryptographic operations of key generation, encryption, decryption, and evaluation. Our experiments show that 80% of computation overhead is attributed to key generation. We next present the costs with and without key generation to study the impact of the member and data size.

Figure 5 (a-b) presents the computation cost with varying number of members. Each members data size used in policy negotiation is set to 5000 data samples. Figure 5 (a) presents the cost of the total computation time excluding key generation. There is a linear increase in time with the growing number of members. This is the fundamental cost of encryption and evaluation that is dominated by matrix encryption and addition. To profile the generation of key cost, in Figure 5 (b), we conducted similar experiments. Each input size cost is increased by the overhead of key generation. The increase is quadratic as a number of slots (plaintext elements) are set to square of input size not to lose any data during input conversion. Yet, the cost is independent of the member size because the key generation is required only once in a computation round. We note that the key generation is not a limiting factor; members may generate keys before a consortium is established.

In Figure 5 (c-d), we show the costs associated with different data samples. The number of members is set to 20. Similar to the previous experiments, the computation costs are dominated by key generation. Our experiments also reported of no effect between the cost and number of samples. As the size of the data samples increases, the overhead is amortized over the operations on local statistics (which is a square matrix of size inputs), and the performance converges that of the input size. This explains the similar trends observed in plots.

To conclude, Curie allows 50 members to be able to compute the secure algorithm with 5K data samples with 40 features in less than a minute. Such experiments above can be used to do a trade-off between data exchange policy and performance. Note that not all changes to a policy specification need to be secure computation-specific. For instance, the use of homomorphic encryption impacts performance but not policy construction and negotiation.

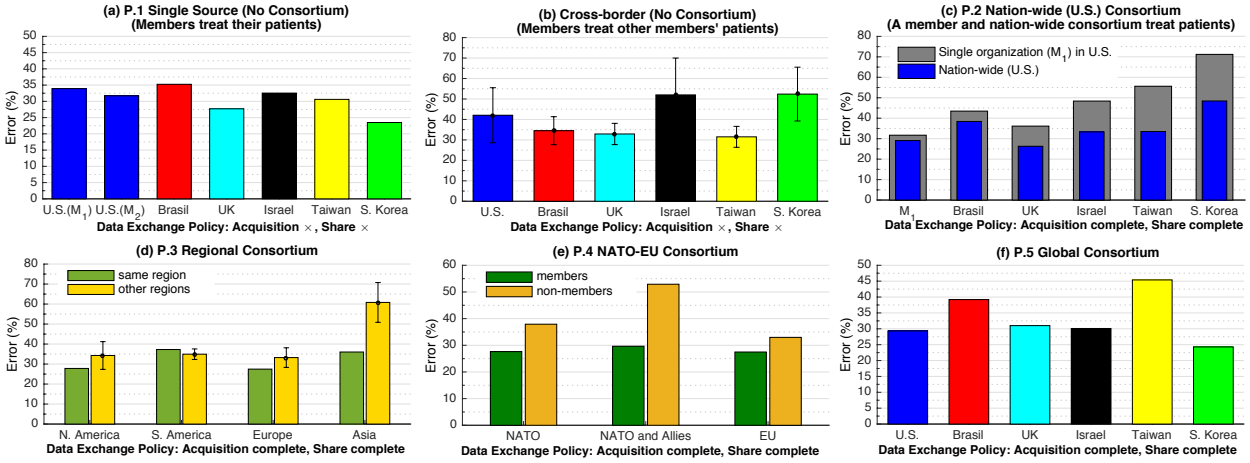
(a) P.1 Single Source (No Consortium)
(Members treat their patients)
Error (%)
U.S.(M1) U.S.(M2) Brasil UK Israel Taiwan S. Korea
Data Exchange Policy: Acquisition ×, Share ×
(b) Cross-border (No Consortium)
(Members treat other members' patients)
Error (%)
U.S. Brasil UK Israel Taiwan S. Korea
Data Exchange Policy: Acquisition ×, Share ×
(c) P.2 Nation-wide (U.S.) Consortium
(A member and nation-wide consortium treat patients)
Single organization (M1) in U.S.
Nation-wide (U.S.)
Error (%)
M1 Brasil UK Israel Taiwan S. Korea
Data Exchange Policy: Acquisition complete, Share complete
(d) P.3 Regional Consortium
same region
other regions
Error (%)
N. America S. America Europe Asia
Data Exchange Policy: Acquisition complete, Share complete
(e) P.4 NATO-EU Consortium
members
non-members
Error (%)
NATO NATO and Allies EU
Data Exchange Policy: Acquisition complete, Share complete
(f) P.5 Global Consortium
Error (%)
U.S. Brasil UK Israel Taiwan S. Korea
Data Exchange Policy: Acquisition complete, Share complete

Figure 6: Implication of policies on accuracy - errors are validated in consortia through data exchange policies.

6.2 Implications of Policy on Results

Data exchange policies allow a member to construct an algorithm beyond those derivable from any single study. The algorithm efficiency should reflect the performance on samples of the members and non-members. This is translated to warfarin study as follows: how well a member of a consortium prescribes warfarin dose for consortia and non-consortia members? and what are the consequences of errors on patient health?. Below, we validate the performance of the algorithm in nation-wide and cross-boundary consortia through data exchange policies.

In our experiments, we used subsets of patient records that had no missing inputs for accurate evaluation. We present the dataset distribution of members used throughout in Appendix C. We split the data into two cohorts; the training cohort is used to learn the algorithm, and the validation cohort is used to assign dose to the new patients based on the data exchange policies. The performance of policies is validated with Mean Absolute Error (MAE) and Mean Absolute Percentage Error (MAPE). These metrics are measures of how far predicted dosages are away from true dose and accuracy as a percentage of the error. The lower values of MAE and MAPE indicate better quality of treatment.

Implications on Accuracy- In our first set of experiments, we validate how well a member prescribe warfarin dose for its local patients and patients traversing borders (*roaming patient*). The results are used as a baseline for comparison of varying consortia and data exchange policies. Figure 6 (a) sought to identify the local algorithm errors (**P.1**)[2]. The errors significantly differ between countries and even for members in the same country (depicted as M_1 and M_2 in U.S.). The lower results are because of having homogeneous data, all the inputs in these countries have similar traits. For instance, similar age, ethnicity, or genetic markers found in a member produce an over-fitted algorithm for its local patients. These findings are validated with use of local algorithms for treatment of other countries' patients. As illustrated in Figure 6 (b) errors yields significantly high for particular countries. The results indicate that improvements for local patients and patients traversing borders lay in creation of policies to increase the patient diversity.

The next experiments measure the error of nation-wide (**P.2**), regional (**P.3**), NATO-EU (**P.4**) and global (**P.5**) policies for the patients of members and non-members. In our experiments, we make three major observations. First, nation-wide policy increases the dose accuracy compared to that of a single member in a country. This result validated through nationwide consortia and a single member (M_1) in United States (See Figure 6 (c)). Second, supporting previous findings, all regional (excluding Asia) and NATO-EU policies decrease the error both for treatment of their patients and other countries patients (See Figure 6 (d-e)). However, Asia collaboration results in unexpected errors, particularly for treatment of other regions' patients. This is because nation-wide, regional and NATO-EU policies include patient population having different characteristics, thus better generalizes to the dosages. In contrast, Asia collaboration lacks large enough White and Black racial groups. Finally, the global collaboration results in higher errors when evaluated for particular countries such as Brazil and Taiwan (See Figure 6 (f)). To conclude, while data exchange policies are effective in reducing errors, the results highlight an important direction of systematic exchange of patient inputs for treating patients based on their race groups.

[2]Our results are consistent with that of findings in reference paper [27]. Pharmacogenetic algorithm is the current best-known algorithm for predicting the initial dose and outperforms clinical and traditional fixed dose approaches as doses can vary by a factor of 10 among individuals.

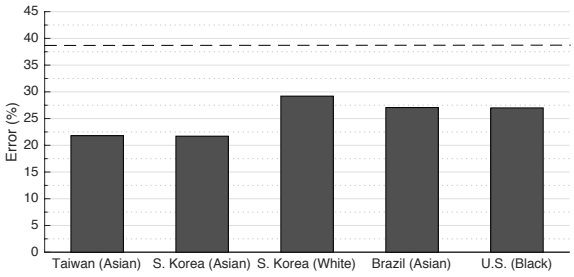

Figure 7: An exploration of impact of data-dependent conditionals and selections defined in global consortium - Members construct an algorithm per race. The dashed line is the average error found without use of conditionals and selections.

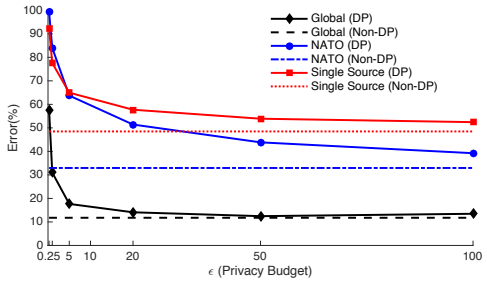

Figure 8: Non-private secure algorithm (Non-DP) vs differentially-private secure algorithm (DP) performance of a member measured against various policies.

To eliminate the errors caused by race differences, members aim at having an ideal population uniformity based on race groups. Presented in Figure 7, the consortia members able to increase the accuracy with the effective use of policies. For instance, a member having a small number of white patients defines selections to acquire solely that group and a member having large enough patients for all races sets data-dependent conditionals to obtain patient inputs that are not similar in its records (e.g., a different genotype). We note that use of different data-dependent conditional algorithms does not result in statistically significant improvements in dataset we use.

We conclude that, based on the above observations, the members are in need for a careful calibration of policies based on the domain- and task-specific properties to maximize the output of the computation.

Consortium	U	SW	O	Selections	Conditionals
Single Source	37.7%	43.4%	18.8%	✗	✗
Nation-wide	18.9%	52.3%	28.8%	✓	✓
NATO	19.3%	51.5%	29.2%	✓	✓
Regional	19%	51.3%	29.7%	✓	✓
Global	21.2%	46.8%	32%	✓	✓

Table 5: Quantifying the impact of policies on health-related risks: The percentage of patients that are over- and -under prescribed and in safety window. Results are from the consortia patients using the algorithm of a member located in the U.S.

Implications on Patient Health- We examine the impacts of the errors found in the previous section on patient health. A study is presented to evaluate the clinical relevance of errors through policies.

Warfarin ranks among top ten drugs in the United States associated with severe or fatal adverse effects caused by incorrect dose treatments [1]. To identify the side effects of warfarin, we follow a medical guide to identify the consequences of under- and over-prescriptions [20]. Taking high-dose warfarin causes thin blood which may result in intracranial and extracranial bleeding. Taking low doses causes thick blood which may result in embolism and stroke. In both cases, extremely high or low doses may lead to higher risks of death.

We define errors that are inside and outside of the warfarin safety window, and the under- or over prescriptions. We consider weekly errors for each patient because using weekly values eliminates the errors posed by the initial (daily) dose. For instance, a small initial dose error might incur more health risks over time. The weekly dose is admitted in the safety window if an estimated dose falls within 20% of its corresponding clinically-deduced value. This is because the value defines the difference of 1 mg per day relative to the fixed starting dose of 5 mg per day which is often accepted as clinically relevant [27,28]. The deviations fall outside of the safety window is an under- or over-prescriptions, and causing health related risks.

Table 5 presents the percentage of patients falls in safety window, under- or over-prescriptions with varying policies of a member. We find that use of policies increases the number of patients in a safety window; thus, they are effective at preventing errors that introduce health-related risks.

6.3 Differential Privacy

To protect individual privacy in dose algorithm, each member may compute the differentially-private secure algorithm on the negotiated data (introduced in Section 5.4). This section presents the results of using the differential-private secure algorithm (DP) instead of using secure dose algorithm (Non-DP).

To establish a baseline performance, we constructed non-private secure algorithms of a member. We then build the differential-private secure algorithm for different privacy budgets ($\varepsilon = 0.25, 1, 5, 20, 50$ and 100), and compare the results of two algorithms through varying consortia.

Figure 8 shows the results of a member in U.S. applying the both algorithms to predict dosages. The algorithms are constructed for the single source, NATO, and global

consortia. The policies are set to complete share and acquisition (i.e., without dictating selections and conditionals). The average error over 100 distinct models for each parameter is reported. We note that other consortia and policies show the similar effect on the results.

Although not surprising, it degrades the accuracy as the ε value increases. For instance, the accuracy improvement obtained through NATO policy over single source degrades with the privacy budget less than or equal to 20. To conclude, the benefit obtained from non-private secure algorithm through data exchange policies vanishes with lower privacy budgets, and hence significantly affects the errors on dose prescriptions and will cause health-related risks as our previous analysis illustrated.

7 Discussion

The preceding analysis of Curie showed that it can be used by members to dictate their data exchange policies.

One requirement for interpreting the policies correctly is a shared schema to solve the compatibility issues among members. For instance, members may interpret the data columns (e.g., column names and types) differently or may not store some information about consortia (e.g., membership status of an alliance). CPL implements a shared schema describing column names, their types, and explanations of data fields as well as consortium-specific information. Members negotiate the schema similar to the policy negotiations and revise the schema based on the schema of the negotiation requester.

In Curie policy language, a set of predefined statistical functions (e.g., intersection size) is provided to securely evaluate over members' data. However, there are other functions suitable that will produce different criteria for conditionals to make decisions about share and acquisition policies and can be pertinent in particular domains. For instance, data exchange among finance companies may require functions to calculate a particular similarity between data distributions. Curie allows easy integration of such functions with conditionals. Future work should investigate the use of different statistical functions.

Finally, we would like to point out that we did not focus much on the reasons of policy impacts on the predictive success of models and its adverse outcomes over time (e.g., patient health). However, our evaluation results showed that CPL policies are able to express both complex relations and address complex constraints on sharing. Hence, it allows members to establish true partnerships and systematic data exchange. While this explanation matches both our intuition and the experimental results, a further domain-specific formal analysis is needed. We plan to pursue this in future work.

8 Related Work

Policy has been used in several contexts as a vehicle for representing configuration of secure groups [31], network management [40], threat mitigation [21, 22], access control [16], patient-driven privacy control [8], data retrievel systems [19] and privileged-augmented detection [9]. These approaches define a schema for their target problem and mostly consider all or nothing proposition while sharing the data or establishing the partnerships. In contrast, Curie is a novel approach in which a policy is used to regulate the data exchange requirements of members that has complex relationships.

On the other hand, secure computation on sensitive proprietary data has recently attracted attention. Anonymization [18], multi-site statistical models [14], and secure multi-party computation [29] have started to shed light on this issue. Such techniques have been used both for training and classification phases in deep learning [41], clustering [25], and decision trees [6]. To allow programmers to develop such applications, efficient secure computation programming frameworks are designed for general purposes [4, 30] and for specific domains such as healthcare [11]. However, these approaches do not consider complex relationships among members and assume all or nothing proposition for exchanging data. We view our efforts in this paper to be complementary to much of these works. They can be easily integrated into Curie to establish partnerships and manage data exchange policies before a computation starts.

9 Conclusion

In this paper, we presented an approach called Curie that allows members to exchange data in complex consortia. Curie provides a policy language, called CPL, to specify the data exchange requirements. The requirements are implemented through share and acquisition policies of members that allow members to dictate conditionals and selections over their data. These policies specify both the restrictions on data exchange and the demand for others' data. We showed that negotiations of pairwise policies are tractable. We validated the use of Curie through varying policies of consortia in an example health care application. A multi-party data exchange algorithm was compiled with the use of secure multi-party computations enhanced with differential privacy illustrating the policy/performance trade-offs. Curie incurs low computation overhead and allows members to establish true partnerships and dictate the requirements for data exchange. Further, we observed that use of policies improves the accuracy and in the long term increases the positive outcomes.

Future work will investigate the use of Curie in other domains focusing on privacy-preserving algorithms and

exploring different metrics for data-dependent conditionals that measure other criteria of exchanging data. We will also integrate Curie into off-the-shelf frameworks designed for secure computation to evaluate its policy/performance trade-offs.

Acknowledgment

We thank Dr. David Lopez-Paz (Facebook AI Research) for the insights in evaluation of warfarin experiments, Dr. Jalaj Upadhyay (post-doctoral fellow at the Penn State University) for suggestions in secure and differentially-private computations on negotiated data and Nicolas Papernot (PhD. student at the Penn State University) for feedback on early versions of this document. Research was sponsored by the Army Research Laboratory and was accomplished under Cooperative Agreement Number W911NF-13-2-0045 (ARL Cyber Security CRA). The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Army Research Laboratory or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation here on.

References

[1] ABRAHAM, N. S., ET AL. Comparative risk of gastrointestinal bleeding with dabigatran, rivaroxaban, and warfarin: population based cohort study. *British Medical Journal* (2015).

[2] BEN-DAVID, A., NISAN, N., AND PINKAS, B. Fairplaymp: a system for secure multi-party computation. In *Computer and communications security* (2008).

[3] BLUNDO, C., DE CRISTOFARO, E., AND GASTI, P. Espresso: efficient privacy-preserving evaluation of sample set similarity. In *Data Privacy Management and Autonomous Spontaneous Security* (2013).

[4] BOGDANOV, D., KAMM, L., LAUR, S., AND SOKK, V. Rmind: a tool for cryptographically secure statistical analysis. *IEEE Transactions on Dependable and Secure Computing* (2016).

[5] BOGDANOV, D., LAUR, S., AND WILLEMSON, J. Sharemind: A framework for fast privacy-preserving computations. In *European Symposium on Research in Computer Security* (2008).

[6] BOST, R., POPA, R. A., TU, S., AND GOLDWASSER, S. Machine learning classification over encrypted data. In *NDSS* (2015).

[7] BRAKERSKI, Z., GENTRY, C., AND VAIKUNTANATHAN, V. (leveled) fully homomorphic encryption without bootstrapping. In *Theoretical Computer Science* (2012).

[8] CELIK, Z. B., LOPEZ-PAZ, D., AND MCDANIEL, P. Patient-driven privacy control through generalized distillation, 2016.

[9] CELIK, Z. B., MCDANIEL, P., IZMAILOV, R., PAPERNOT, N., AND SWAMI, A. Extending detection with forensic information, 2016.

[10] CHAUM, D., CRÉPEAU, C., AND DAMGARD, I. Multiparty unconditionally secure protocols. In *Theory of computing* (1988).

[11] CHIDA, K., MOROHASHI, G., FUJI, H., MAGATA, F., FUJIMURA, A., HAMADA, K., IKARASHI, D., AND YAMAMOTO, R. Implementation and evaluation of an efficient secure computation system using r for healthcare statistics. *Journal of the American Medical Informatics Association* (2014).

[12] DAMGÅRD, I., KELLER, M., LARRAIA, E., PASTRO, V., SCHOLL, P., AND SMART, N. P. Practical covertly secure mpc for dishonest majority–or: Breaking the spdz limits. In *European Symposium on Research in Computer Security* (2013).

[13] DAMGÅRD, I., PASTRO, V., SMART, N., AND ZAKARIAS, S. Multiparty computation from somewhat homomorphic encryption. In *CRYPTO*. 2012.

[14] DANKAR, F. K. Privacy preserving linear regression on distributed databases. *Transactions on Data Privacy* (2015).

[15] DE CRISTOFARO, E., GASTI, P., AND TSUDIK, G. Fast and private computation of cardinality of set intersection and union. In *Cryptology and Network Security* (2012).

[16] DUAN, L., ZHANG, Y., CHEN, S., ZHAO, S., WANG, S., LIU, D., LIU, R. P., CHENG, B., AND CHEN, J. Automated policy combination for secure data sharing in cross-organizational collaborations. *IEEE Access* (2016).

[17] DWORK, C. Differential privacy. In *Encyclopedia of Cryptography and Security* (2011).

[18] EL EMAM, K., SAMET, S., ARBUCKLE, L., TAMBLYN, R., EARLE, C., AND KANTARCIOGLU, M. A secure distributed logistic regression protocol for the detection of rare adverse drug events. *Journal of the American Medical Informatics Association* (2013).

[19] ELNIKETY, E., MEHTA, A., VAHLDIEK-OBERWAGNER, A., GARG, D., AND DRUSCHEL, P. Thoth: Comprehensive policy compliance in data retrieval systems. In *USENIX Security* (2016).

[20] Medication guide, caumadin (warfarin sodium). http://www.fda.gov. [Online; accessed 11-July-2016].

[21] FREUDIGER, J., DE CRISTOFARO, E., AND BRITO, A. E. Controlled data sharing for collaborative predictive blacklisting. In *Proc. DIMVA* (2015).

[22] GARRIDO-PELAZ, R., GONZÁLEZ-MANZANO, L., AND PASTRANA, S. Shall we collaborate?: A model to analyse the benefits of information sharing. In *Proceedings of the 2016 ACM on Workshop on Information Sharing and Collaborative Security* (2016), ACM, pp. 15–24.

[23] GENTRY, C. *A fully homomorphic encryption scheme*. PhD thesis, Stanford University, 2009.

[24] GOLDREICH, O. *Foundations of cryptography: volume 2, basic applications*. Cambridge university press, 2009.

[25] GRAEPEL, T., LAUTER, K., AND NAEHRIG, M. Ml confidential: Machine learning on encrypted data. In *Information Security and Cryptology* (2012).

[26] HUANG, Y., EVANS, D., KATZ, J., AND MALKA, L. Faster secure two-party computation using garbled circuits. In *USENIX Security Symposium* (2011).

[27] INTERNATIONAL WARFARIN PHARMACOGENETICS CONSORTIUM. Estimation of the warfarin dose with clinical and pharmacogenetic data. *The New England journal of medicine* (2009).

[28] KIMMEL, S. E., ET AL. A pharmacogenetic versus a clinical algorithm for warfarin dosing. *New England Journal of Medicine* (2013).

[29] LINDELL, Y., AND PINKAS, B. Secure multiparty computation for privacypreserving data mining. *Journal of Privacy and Confidentiality* (2009).

[30] LIU, C., WANG, X. S., NAYAK, K., HUANG, Y., AND SHI, E. Oblivm: A programming framework for secure computation. In *Security and Privacy* (2015).

[31] MCDANIEL, P., AND PRAKASH, A. Methods and limitations of security policy reconciliation. *ACM Transactions on Information and System Security (TISSEC)* (2006).

[32] MCDANIEL, P., PRAKASH, A., AND HONEYMAN, P. Antigone: A flexible framework for secure group communication. In *USENIX Security* (1999), pp. 99–114.

[33] MCGREGOR, A., MIRONOV, I., PITASSI, T., REINGOLD, O., TALWAR, K., AND VADHAN, S. The limits of two-party differential privacy. In *Foundations of Computer Science (FOCS)* (2010).

[34] NIELSEN, J. B., NORDHOLT, P. S., ORLANDI, C., AND BURRA, S. S. A new approach to practical active-secure two-party computation. In *CRYPTO*. 2012.

[35] American recovery and reinvestment act of 2009. https://en.wikipedia.org/wiki/American_Recovery_and_Reinvestment_Act_of_2009. [Online; accessed 01-December-2016].

[36] Health information technology for economic and clinical health act. https://en.wikipedia.org/wiki/Health_Information_Technology_for_Economic_and_Clinical_Health_Act. [Online; accessed 01-December-2016].

[37] Nato standard ajp-4.10 allied joint doctrine for medical support. http://www.nato.int/cps/en/natohq/topics_49168.htm? [Online; accessed 01-December-2016].

[38] REPORT, E. C. An implementation of homomorphic encryption. https://github.com/shaih/HElib. [Online; accessed 01-January-2017].

[39] REPORT, E. C. Overview of the national laws on electronic health records in the eu member states and their interaction with the provision of cross-border ehealth services. http://ec.europa.eu/index_en.htm. [Online; accessed 01-December-2016].

[40] RIEKSTIN, A. C., JANUÁRIO, G. C., RODRIGUES, B. B., NASCIMENTO, V. T., CARVALHO, T. C., AND MEIROSU, C. Orchestration of energy efficiency capabilities in networks. *Journal of Network and Computer Applications* (2016).

[41] SHOKRI, R., AND SHMATIKOV, V. Privacy-preserving deep learning. In *Proc. ACM CCS* (2015).

[42] SOLOVE, D. J., AND SCHWARTZ, P. M. *Information privacy law*. Aspen Publishers, 2015.

[43] WU, F.-J., KAO, Y.-F., AND TSENG, Y.-C. From wireless sensor networks towards cyber physical systems. *Pervasive and Mobile Computing* (2011).

[44] WU, G., HE, Y., WU, J., AND XIA, X. Inherit differential privacy in distributed setting: Multiparty randomized function computation. *arXiv preprint arXiv:1604.03001* (2016).

[45] WU, X., FREDRIKSON, M., WU, W., JHA, S., AND NAUGHTON, J. F. Revisiting differentially private regression: Lessons from learning theory and their consequences. *arXiv preprint arXiv:1512.06388* (2015).

[46] YAO, A. C. Protocols for secure computations. In *Foundations of Computer Science* (1982).

[47] ZHANG, J., ZHANG, Z., XIAO, X., YANG, Y., AND WINSLETT, M. Functional mechanism: regression analysis under differential privacy. *VLDB* (2012).

A Curie Policy Language

This section presents the Backus Naur Form (BNF) of Curie data exchange policy language.

⟨curie_policy⟩ ::= ⟨statements⟩

⟨statements⟩ ::= ⟨statement⟩ ';' [⟨statements⟩]

⟨statement⟩ ::= ⟨share_clause⟩
| ⟨acquire_clause⟩
| ⟨attribute⟩
| ⟨sub_clause⟩

; share clauses defined as follows:

⟨share_clause⟩ ::= 'share' ':' [⟨members⟩] ':' [⟨conditionals⟩] '::' ⟨selections⟩

; acquisition clauses defined as follows:

⟨acquire_clause⟩ ::= 'acquire' ':' [⟨members⟩] ':' [⟨conditionals⟩] '::' ⟨selections⟩

; attributes are defined as follows:

⟨attribute⟩ ::= ⟨identifier⟩ ‘:=’ ‘<’ ⟨value⟩ ‘>’
| ⟨identifier⟩ ‘:=’ ‘<’ ⟨value_list⟩ ‘>’

; user defined sub-clauses defined as follows:

⟨sub_clause⟩ ::= ⟨tag⟩ ‘:’ [⟨conditionals⟩] ‘::’ ⟨selections⟩

; conditionals including data-dependent function evaluations defined as follows:

⟨conditionals⟩ ::= ⟨var⟩‘=’⟨value⟩ [‘,’ ⟨conditionals⟩]
| ‘evaluate’ ‘(’ ⟨data_ref⟩ ‘,’ ⟨alg_arg⟩ ‘,’ ⟨threshold_arg⟩ ‘)’ [‘,’ ⟨conditionals⟩] | ‘’

⟨selections⟩ ::= ⟨filters⟩
| ⟨tag⟩

⟨filters⟩ ::= ⟨filter⟩ [‘,’ ⟨filters⟩]

⟨filter⟩ ::= ⟨var⟩ ⟨operation⟩ ⟨value⟩ | ‘’

⟨data_ref⟩ ::= ‘&’ ⟨identifier⟩

⟨alg_arg⟩ ::= ⟨algorithms⟩

⟨algorithms⟩ ::= ‘Intersection size’
| ‘Jaccard index’
| ‘Pearson correlation’
| ‘Cosine similarity’

⟨threshold_arg⟩ ::= ⟨floating_point_number⟩

⟨operation⟩ ::= ‘=’ | ‘<’ | ‘>’ | ‘!=’ | in |

⟨value_list⟩ ::= ‘{’ ⟨value⟩ ‘}’ [‘,’ ⟨value_list⟩]

⟨members⟩ ::= ⟨member⟩ [‘,’ ⟨members⟩]

⟨member⟩ ::= ⟨identifier⟩ | ‘’

; for completeness, trivial items defined as follows:

⟨identifier⟩ ::= ⟨word⟩

⟨var⟩ ::= ‘$’ ⟨identifier⟩

⟨value⟩ ::= ⟨string⟩

⟨tag⟩ ::= ⟨word⟩

⟨string⟩ ::= ‘"’ ⟨stringchars⟩ ‘"’

⟨stringchars⟩ ::= ⟨stringletter⟩ [⟨stringchars⟩]

⟨stringletter⟩ ::= 0x10 | 0x13|0x20| ... | 0x7F

⟨word⟩ ::= ⟨char⟩ [⟨word⟩]

⟨char⟩ ::= ⟨letter⟩ | ⟨digit⟩

⟨letter⟩ ::= ’A’ | ’B’ | ... | ’Z’ | ’a’ | ’b’ | ... | ’z’ | 0x80 | 0x81 | ... | 0xFF

⟨digit⟩ ::= ’0’ | ’1’ | ... | ’9’

⟨floating_point_number⟩ ::= ⟨decimal_number⟩ ’.’ [⟨decimal_number⟩]

⟨decimal_number⟩ ::= ⟨digit⟩ [⟨decimal_number⟩]

B Members of Data Collection

Our study uses a data set collected by the International Warfarin Pharmacogenetics Consortium (IWPC), to date the most comprehensive database containing patient data collected from 24 medical institutions from 9 countries [27]. Figure 2 shows the medical institutions involved in the study. The dataset does not include the name of the medical institutions, yet there is a separate ethnicity dataset provided for identifying the genomic impacts of the algorithm. We use the race (reported by patients) and race (categories defined by the Office of Management and Budget) to identify the country of a patient[3]. For instance, we consider a medical institution with a high number of Japanese race is located in Japan.

C Data Distribution

Figure 1 illustrates the exploration of local patient data distribution used in policy evaluation in Section 6.2.

D Secure Computation Analysis

In this appendix, we present accuracy and privacy guarantees of dose algorithm depicted in Figure 2. Following our discussion in Section 5.3, secure computation layer provides an accurate and private algorithm for all parties (i.e., members) through share of encrypted integrated statistics ($O_i = X^\intercal X$ and $\mathcal{V}_i = X^\intercal Y$ matrices). Since all data exchange is encrypted through the homomorphic encryption, the security of the algorithm against any adversary outside the authorized parties is based on the underlying homomorphic cryptosystem. In the following, we analyze the privacy guarantees in the presence of two different semi-honest adversary models[4] according to the distribution of the corrupted parties[5]:

An adversary “not” involving session initiator- Assume for now that a session initiator is trusted and does not collude with other parties. Loosely speaking, since all computations are performed on the encrypted data, none of the parties learn anything about other parties’ input. To prove, we define an m-party protocol $\mathcal{P}$ to compute a function f where $f : (x_1, \ldots, x_m) \rightarrow (y_1, \ldots, y_m)$ is an m-ary functionality [24]. We donate the view of i^{th} party during execution of the protocol on inputs $(x_1, \ldots, x_n)$ as $View_{\mathcal{P}_i}(\bar{x})$. The view includes input x_i, random string r_i, and all received messages such that it is defined as $(x_i, r_i, m_{i_1}, \ldots, m_{i_t})$ after t rounds of the protocol. The

[3]The authors indicated via personal communication that they cannot provide the exact name of the institutions due to the privacy concerns.
[4]Same as the honest-but-curious adversary model
[5]We only consider the non-adaptive adversary model and modifying an adversary’s its input before joining the session is not considered as a security breach.

output of parties during the execution of a view the becomes $Output_{\mathcal{P}_l}(\bar{x})$ for the rounds $I = i_1, \ldots, i_t \subseteq [m] \stackrel{def}{=} \{1, \ldots, m\}$ and the inputs $\bar{x} = (x_1, \ldots, x_n)$. In this setting, the protocol privately computes f in the presence of semi-honest parties if there exists a probabilistic polynomial-time simulator S for every message $I \subseteq [M]$ such that

$$\Big(S(I, (x_{i_1}, ..., x_{i_t}), f(\bar{x}))\Big) \stackrel{C}{\equiv} \Big(View_{\mathcal{P}_l}(\bar{x}), Output_{\mathcal{P}_l}(\bar{x})\Big) \tag{1}$$

where $\stackrel{C}{\equiv}$ shows computational indistinguishability.

To explain the intuition behind the equation 1, we first consider that one of the parties is semi-honest. If a protocol $\mathcal{P}$ can securely compute a function f in the presence of that corrupted party, then the view of that party can be simulated by a probabilistic polynomial-time simulator using only inputs and outputs of this party. Simply put, a corrupted party learns from the execution of the protocol is already implied by its output but not more. Equation 1 indicates that this security guarantee should hold for any subset of m-parties. Before starting to analyze our secure dose algorithm, we note that the output of it will be revealed to one party (i.e., session initiator) only in each round and there is no any dependency between the rounds. Hence, below, we only show the guarantees for one round, but they are held for every round of the protocol.

We consider the party P_{i+1} in Figure 2, the view of party P_{i+1} consists of the public key generated by the session initiator K, the encryption of local statistics of previous parties $M_i = < E(O_i)_K, E(\mathcal{V}_i)_K >$. Its input is $< \mathcal{V}_{i+1}, O_{i+1} >$ and its output is $M_{i+1} = < E(O_i + O_{i+1}), E(\mathcal{V}_i + \mathcal{V}_{i+1}) >$. A simulator S selects random values for its own inputs $< \mathcal{V}''_{i+1}, O'_{i+1} >$ and encrypts them using the public key published by the session initiator. Then, the simulator S performs the homomorphic operation on the received message M_i and outputs $M'_{i+1} = < E(O_i + O'_{i+1})_K, E(\mathcal{V}_i + \mathcal{V}''_{i+1})_K >$. At this point, we assume the underlying homomorphic encryption is semantically secure, i.e., it is infeasible for a computationally bounded adversary to derive significant information about the corresponding plaintext when given only its ciphertext and the public key. Therefore, the output of the simulator M'_{i+1} is computationally indistinguishable from output of the real execution of the protocol M_{i+1} for every input pairs. Therefore, the protocol privately computes the function in the presence of one semi-honest corrupted party.

The extension to multi-corrupted semi-honest adversaries is straightforward as the only difference is the view of a subset of parties having many encrypted messages. Since the semantic security of the underlying homomorphic encryption is hold for any pair of these many encrypted messages, no information leaks about the corresponding plaintexts. Therefore, parties privately compute the function f in the presence of semi-honest parties.

Adversary involving session initiator- We consider the case when session initiator is corrupted. In this setting, session initiator has the session private key; it can decrypt the encrypted data using its private key. Therefore, the encryption layer does not provide security guarantees. The corrupted parties including session initiator can infer the input of an honest party if the predecessor (previous party) and successor (next party) of an honest party are both corrupted. We consider the possible cases for the leakage as follows:

1. 2-party: The session initiator is corrupted, and another party is honest. In this case, predecessor and successor of the honest party are both the corrupted session initiator. Therefore, the input of honest party is learned by the corrupted party.
2. 3-party: A corrupted session initiator is either predecessor or successor, and it can learn one honest parties input only if another party is corrupted.
3. n-party ($n > 3$): Similar to the previous cases, the probability of the input leakage depends on the placement of the corrupted parties. To learn an honest party's input, at least two parties must be corrupted and placed in previous and next of the honest party.

Although the individual raw data itself does not leak, the risk of inappropriate disclosures from local summary statistics exists in some extreme cases [18]. For example, consider the exchange of plain matrix $V_i = X^\intercal Y$ among two parties, a party may use the extreme values found in V_i to identify particular patients. For instance, in dose algorithm, taking inducers such as Rifadin and Dilantin could indicate high dose prescriptions. If the values in V_i are high, then a party may infer patients taking enzyme inducers and the presence of high dosage warfarin intake (which suggests the presence of patients with thick blood). Similarly, exchange of $O_i = X^\intercal X$ may leak information about the number of observations and represent the number of Os or 1s in a column. For instance, for the former first entry in the matrix, $X^\intercal X$, gives the total number of patients. For the latter, $(X^\intercal X)_{j,j}$ gives the number of 1s in the column. This type information lets a party infer knowledge, particularly when binary inputs (e.g., use of medicine) are used.

To summarize, we showed that our protocol privately computes the secure dose algorithm against an adversary which is not involving the session initiator. The policy-based approach supports this assumption, as parties decide their local policies according to regulations and laws.

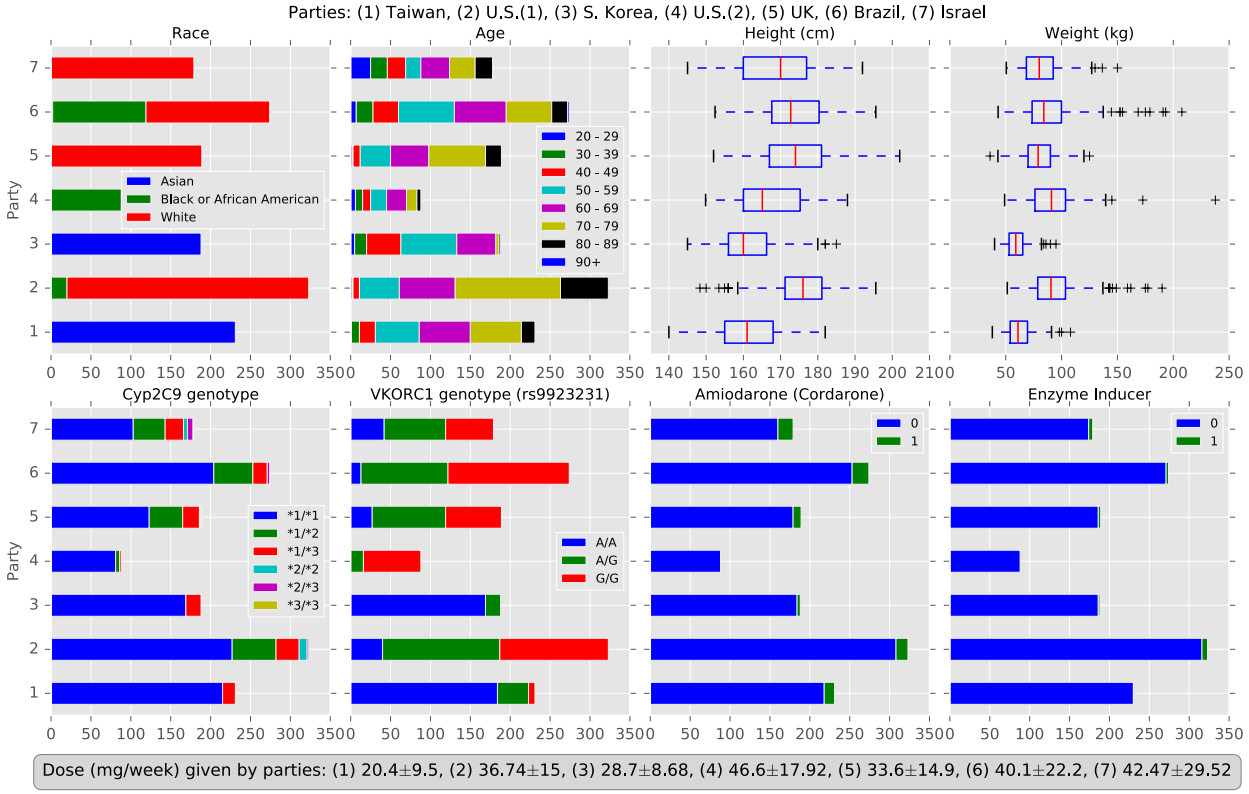

Figure 1: An exploration of patient inputs. We illustrate the major medical institutions involved in data collection that have no missing patient inputs. (We suggest printing in color for better visualizations.)

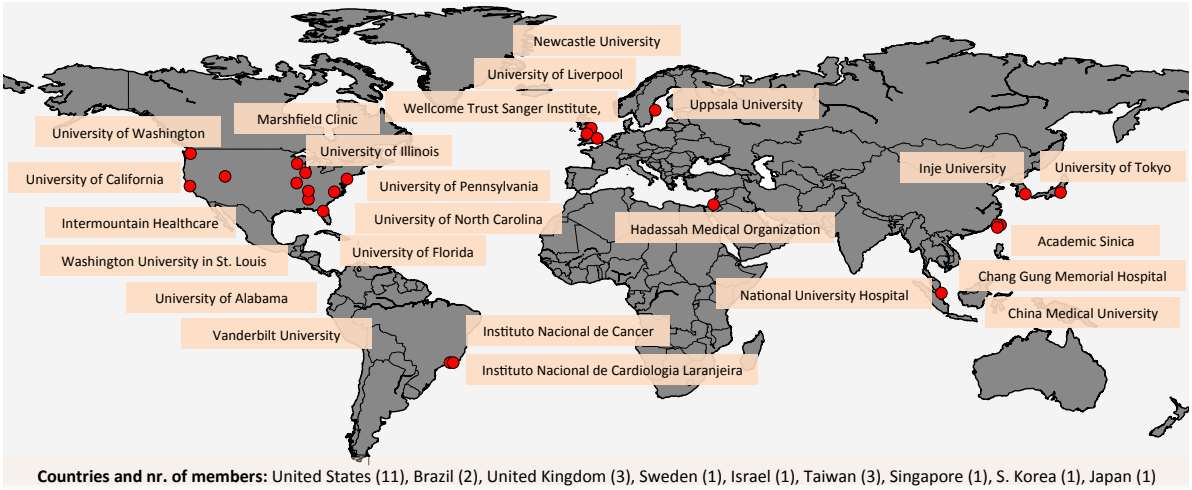
Newcastle University
University of Liverpool
Wellcome Trust Sanger Institute,
Uppsala University
Marshfield Clinic
University of Washington
University of Illinois
University of California
University of Pennsylvania
Inje University
University of Tokyo
Intermountain Healthcare
University of North Carolina
Hadassah Medical Organization
Academic Sinica
Washington University in St. Louis
University of Florida
Chang Gung Memorial Hospital
University of Alabama
National University Hospital
China Medical University
Vanderbilt University
Instituto Nacional de Cancer
Instituto Nacional de Cardiologia Laranjeira
Countries and nr. of members: United States (11), Brazil (2), United Kingdom (3), Sweden (1), Israel (1), Taiwan (3), Singapore (1), S. Korea (1), Japan (1)

Figure 2: Medical institutions from nine countries and four continents involved in data collection for warfarin dose algorithm.